

AUTOMATIC CONSTRUCTION OF ACCURATE MODELS OF PHYSICAL SYSTEMS

Final Report

18 August 1999

Office of Naval Research Contract #:

N00014-96-1-0720

End Date:

30 June 1999

Principal Investigator:

Elizabeth Bradley

Address:

University of Colorado

Boulder, CO 80309-0430

lizb@cs.colorado.edu

(303) 492-5355

DISTRIBUTION STATEMENT A

Approved for Public Release

Distribution Unlimited

Summary

System identification (SID) is the task of deducing the internal dynamics of a black-box system solely from observations of its outputs. This is an essential first step in a variety of engineering problems; most traditional control-theoretic methods, for example, require an accurate ordinary differential equation (ODE) model. Accuracy is not the only requirement, however; for efficiency reasons, engineers work hard to construct *minimal* models — ODEs that ignore unimportant detail and capture only the behavior that is important for the task at hand. This tandem of accuracy and abstraction is a subtle and difficult part of an engineer's training and practice.

The computer program PRET, the topic of this report, automates the process described in the previous paragraph by building an artificial intelligence (AI) layer around a set of traditional system identification techniques. This AI layer executes many of the high-level parts of the SID process that are normally performed by a human expert, intelligently assessing the task at hand and then reasoning from that information to automatically choose, invoke, and interpret the results of appropriate lower-level techniques. These tactics guide PRET quickly and accurately to the minimal ODE that accounts for the important behavior of the system.

When this contract began in April 1996, PRET incorporated only a few rules and had only been tested on two simple mechanical systems. It relied on an expert user for observations, and had only been invoked by its primary author. As a result of the work funded by this contract, PRET can now build models of graduate textbook level systems in several domains (mechanical, viscoelastic, electronic), and has successfully solved a challenging real-world engineering problem. It can interact with the physical world automatically, using sensors and actuators to plan, perform, and interpret the results of experiments, and it can be used by practicing engineers who know nothing about artificial intelligence.

19990907 133

DTIC QUALITY INSPECTED 4

The ultimate goal of the PRET project is to build a tool that can automatically construct models of high-dimensional, nonlinear, black-box systems, drawn from any domain that admits ODE models. Our objective is to effectively automate existing SID techniques, not invent new ones; even so, this is an extremely ambitious goal, as nonlinear system identification is a state-of-the-art research area. As an automated tool that accomplishes this, PRET is of obvious practical importance in science and engineering: as a corroborator of existing models and designs, as a medium within which to instruct newcomers, and as an intelligent assistant, whose aid can allow more time and thought to be devoted to other demanding tasks.

FY99 Progress

Our effort during FY99 focused on the *intelligent sensor analysis and actuator control* (ISAAC) subsystem, which lets PRET explore a target system's behavior interactively, using sensors and actuators to perform experiments whose results augment its knowledge in a manner that is useful to the task of modeling that system. Automating the I/O analysis process is hard and interesting, from both AI and engineering standpoints; planning, executing, and interpreting experiments requires some fairly difficult reasoning about what experiments are possible and useful. Distilling available sensor information into qualitative form is comparatively straightforward,¹ but reasoning about the information so derived is surprisingly challenging. Dealing with actuators is even harder; among other things, doing so involves solving control theory's *controllability/reachability* problem: given a system, an initial condition, and an actuator, what state-space points are reachable with the available control input? Finally, in an automated framework, it is also important to reason about the *utility* of a particular experiment. Whether or not PRET can learn anything useful from an experiment depends on what it knows and what it is trying to establish.

As an initial solution to these problems, we have developed two new knowledge representation and reasoning paradigms, *qualitative bifurcation analysis* and *component-based modeling*. The first of these two paradigms is based on a new construct called the *qualitative state/parameter (QS/P) space*, an abstraction of regular state space with an added parameter axis. The QS/P space is a useful way to capture information about actuator signals, sensor data, and different behavioral regions in a single compact representation. The qualitative bifurcation analysis paradigm combines this representation with a set of associated reasoning tools, including geometric reasoning and an extension of a classic reasoning technique from nonlinear dynamics. The second paradigm, *component-based modeling*, exploits generalized physical networks or GPNs, a network-based representation² that makes it easy to incorporate external effects into the model (e.g., loading effects, which can easily be represented directly as part of the network). A set of integrated reasoning techniques exploits these features to help PRET build good models of actuators and actuator/system interactions. The GPN representation has another important feature: it smoothly handles models in multiple domains. This flexibility provides a natural mechanism to adapt the reasoning tools and levels to the amount that PRET knows about the target system, which drastically reduces an otherwise exponential search space.

¹ See our 1998 Intelligent Data Analysis paper on this topic.

² This scheme turned out to have several significant advantages over our initial choice, the bond graph.

We have implemented the skeleton of these two KRR frameworks and reported on the results in refereed conference papers, but we still need to complete the final implementation and test them out on some challenging real-world engineering problems, and we are in the process of seeking funding to do so.

Accomplishments — FY96-99

- PRET efficiently and effectively solves a class of global optimization problems involved in the process of parameter estimation. Specifically, qualitative reasoning (QR) techniques are used to derive starting values for a nonlinear least-squares solver call. These good initial approximations guide the NLS solver away from local extrema in a complicated nonlinear regression landscape. QR techniques are also used to compute cutoff frequencies for digital filters in order to identify and remove noise.

Our initial conference paper on this topic was selected — along with with seven others, out of a total of 76 — for a presentation of the highlights of the 1996 and 1997 QR workshops that was given at IJCAI 1997, and has since been published in the *Annals of Mathematics of Artificial Intelligence*.

- PRET automatically synthesizes model fragments from scratch using power-series techniques. This technique is ubiquitous in engineering, but nonexistent in AI modeling tools. It is not only powerful, but was also very easy to implement: for instance, a one- or two-term expansion in the first few derivatives of a state variable yield every form of force term found in freshman physics textbooks.
- The logic system that forms PRET's core, which uses SLD resolution, static abstraction levels, dynamic meta control, constraint propagation, and a “sparse” truth-maintenance system (STMS) — a significant innovation — to build and check ODE models, is now complete and case-hardened.
- Reinhard Stolle, the author of this logic system, defended his Ph.D., received offers for two tenure-track jobs and a variety of postdoc positions, and published five papers, one in the *Annals of Mathematics of Artificial Intelligence* and one in the highly prestigious National Conference on AI (AAAI). His paper on the dynamic meta control system is in review at *Knowledge and Information Systems*, and a companion paper on the STMS is in preparation. He has been a postdoc at Stanford since January 1999.
- PRET's graphical user interface has been completed and tested.

Apollo Hogan, who designed and implemented that GUI and worked closely with Stolle on the meta control and STMS systems, is a co-author on the *Knowledge and Information Systems* paper listed above. In August 1998, he began a Ph.D. in set-theoretic topology in the Mathematics Department at Berkeley.

- PRET uses simple computer vision techniques to process sensor data and construct high-level observations of a target system. This involved solving part of control theory's *observer problem*: inferring internal system dynamics from incomplete output sensor data. We accomplish this using delay-coordinate embedding and then perform feature extraction on the resulting reconstructed dynamics with some simple geometric reasoning/computer vision techniques.

Our initial conference paper on these results, which was presented at IDA97 in London, was judged to be among the best papers of that conference and was published in a "best of the conference" issue of the *Intelligent Data Analysis* journal.

- [FY99] We have completed the initial design of PRET's intelligent sensor analysis and actuator control (ISAAC) subsystem, which interacts automatically with the target system. This involved the development of two new knowledge representation and reasoning frameworks — qualitative bifurcation analysis and component-based modeling.
- [FY99] Matt Easley, the author of the ISAAC subsystem, defended his Ph.D. thesis proposal and presented three refereed papers, one at the highly competitive International Joint Conference on AI (IJCAI).
- [FY99] Facilitated by these KRR frameworks, we have worked out some ways to improve the model generation process and the automatic planning, execution, and interpretation of experiments, along with some preliminary ideas about the syntax for representation of actuator-related information.
- [FY99] These KRR frameworks also made it almost transparently simple to implement new domains, such as viscoelastics, which has been the focus of state-of-the-art AI modeling programs by other groups. This flexibility and extensibility — and the comparisons to other work that the new domain implementation allowed — proved positive and satisfying.
- [FY99] We have built some simple mechanical and mechatronic test systems with which to test these ideas.
- We have successfully used PRET on a real-world example — a radio-controlled car being used in a robotics system at the University of British Columbia. PRET's model was actually *better* than the one constructed by the project analyst in several senses, and the discrepancies helped the analyst refine his mental model of the system.
- [FY99] Last but not least, the PI was awarded tenure, promoted to Associate Professor, and won the 1999 College of Engineering teaching award.

Transitions

In FY96, we began by verifying PRET's performance on some of the examples used in the graduate system identification course offered by the Aerospace Engineering Department at the University of Colorado. Since then, working with domain experts outside our group, we have successfully used PRET on some real-world problems, including a radio-controlled car being used in a robotics system at the University of British Columbia. The R/C car model that PRET produced was technically correct, but very different from the one constructed by the project analyst, and the disparities surprised him in interesting and useful ways. Specifically, the car's initial velocity was modeled as negative, but the analyst knew that the car had started from rest — a fact that he *had not mentioned in the input information*. This violation of his intuition not only drew this piece of implicit knowledge (" $v = 0$ at $t = 0$ ") out into the explicit syntax of the PRET call, but also suggested a piece of physics that he had not noticed until then — there was a delay in the system. From a cognitive science point of view, this interchange was particularly interesting, as the discrepancies helped the analyst refine his mental model of the system.

The goal of this first round of modeling exercises was simply to duplicate models found by a human analyst, but the R/C car example indicated that PRET had already moved beyond the verification stage — that it can indeed make useful contributions to the modeling art. In order to extend these results, we have initiated collaborations with four experts in a variety of problem domains:

- A mechanical engineer at MIT, Dave Trumper, who has a research project that involves modeling an eyeglass lens cutting machine.
- A physicist at the University of Colorado at Denver, Randy Tagg, who is interested in using PRET to model a ship-mounted crane.
- A hydrologist at the University of Colorado at Boulder, Matt Rutherford, who is interested in using PRET to model water resource systems in the Colorado mountains, for the purposes of decision support in irrigation and flood management.
- A biologist at the University of Colorado at Boulder, Geetha Krishnan, who is interested in using PRET to model hearing loss in the cochlea, a critical step in prosthetic design.

Real-world problems like this have been a rich source of useful insights and suggestions, and PRET's structure and function make this kind of collaboration simply a matter of a few in-person meetings and a dozen or so email messages. Evaluating the results is straightforward; if a controller designed around each of PRET's models succeeds in directing the corresponding application to perform some set of useful actions prescribed by the domain expert, we judge the results successful. The crane, in particular, has obvious and significant naval applications — positioning and stabilizing a load swinging from a deck-mounted crane in rough seas is a difficult control problem, and an accurate ODE model of the ship-crane-load system is essential to its solution.

Finally, we have also collaborated with the Information Technology group at NIST/NOAA in Boulder in order to develop an AI-adapted global optimization tool that combines qualitative reasoning and local numerical methods; see the *Annals of Mathematics and Artificial Intelligence* paper listed below.

Related Publications

FY99:

- E. Bradley, "Time-Series Analysis," in M. Berthold and D. Hand, editors, *Intelligent Data Analysis: An Introduction*, Springer Verlag, 1999.
- M. Easley and E. Bradley, "Generalized physical networks for automated model building," *IJCAI-99 (International Joint Conference on Artificial Intelligence)*, Stockholm; August 1999.
- M. Easley and E. Bradley, "Reasoning About Input-Output Modeling of Dynamical Systems," *IDA-99 (International Workshop on Intelligent Data Analysis)*, Amsterdam; August 1999.
- M. Easley and E. Bradley, "Hybrid phase-portrait analysis in automated system identification," *AAAI Spring Symposium on Hybrid Systems in AI*, Stanford; March 1999.
- Hogan, R. Stolle and E. Bradley, "Putting Declarative Control to Work," in review, *IEEE Transactions on Systems, Man, and Cybernetics*; also available as Technical Report CU-CS (Department of Computer Science) 856-98

FY98:

- E. Bradley, A. O'Gallagher, and J. Rogers, "Global Solutions for Nonlinear Systems using Qualitative Reasoning," *Annals of Mathematics and Artificial Intelligence*, **23**:211-228 (1998)
- E. Bradley and M. Easley, "Reasoning About Sensor Data for Automated System Identification," *Intelligent Data Analysis: An International Journal*, volume 2, number 2, Elsevier Science (1998) [Published electronically and archived on CD; neither format includes page numbers. See www.elsevier.com/locate/ida]
- R. Stolle and E. Bradley, "Multimodal Reasoning for Automatic Model Construction," *Proceedings of the AAAI (National Conference on Artificial Intelligence)*, Madison WI; July 1998

- R. Stolle and E. Bradley, "Multimodal Reasoning about Physical Systems," *Proceedings of the AAAI Spring Symposium on Multimodal Reasoning*; Stanford CA; March 1998. AAAI Technical Report SS-98-04.

FY97:

- E. Bradley and M. Easley, "Reasoning About Sensor Data for Automated System Identification," *Proceedings of the International Workshop on Intelligent Data Analysis*, London UK; August 1997
- E. Bradley, A. O'Gallagher, and J. Rogers, "Global Solutions for Nonlinear Systems using Qualitative Reasoning," *Proceedings of the International Workshop on Qualitative Reasoning about Physical Systems*, Cortona Italy; May 1997
- R. Stolle and E. Bradley, "Opportunistic modeling," *Proceedings of the IJCAI (International Joint Conference on Artificial Intelligence) Workshop on Engineering Problems in Qualitative Reasoning*, Nagoya Japan; August 1997

FY96:

- E. Bradley and R. Stolle, "Automatic Construction of Accurate Models of Physical Systems," *Annals of Mathematics and Artificial Intelligence*, 17:1-28 (1996)
- R. Stolle and E. Bradley, "A Customized Logic Paradigm for Reasoning about Models," *Proceedings of the International Workshop on Qualitative Reasoning about Physical Systems*, Stanford Sierra Camp CA; May 1996